An efficient solution of the pressure Poisson equation using multigrid techniques

Stephan Kramer and Chris Pain

January 28, 2007

1 Introduction

An accurate description of ocean flow is given by the Navier Stokes equations. In the numerical modelling of ocean flow however various approximations must be made to come to realistic computation times. Most ocean models that study global ocean circulation are based on the hydrostatic pressure assumption in which the vertical momentum equation is reduced to a hydrostatic balance between pressure gradient and buoyancy. Although this gives an accurate description of large scale oceanographic phenomena, in many smaller scale problems the vertical acceleration cannot be neglected. Those local nonhydrostatic phenomena however, such as for example the flow over a steep ocean ridge, deep convection, internal wave breaking, etc., may have a large impact on the global ocean circulation.

The computational effort of global ocean modelling with current nonhydrostatic models is generally considered to be too large. An important objective in current numerical oceanography research is therefore to come to ocean models that perform efficiently on the whole range of scales that can be found in the phenomena associated with the global ocean circulation. An important contribution is given in [1]. Their approach is however based on the application of structured grids with no obvious analogue for meshes that are unstructured in all 3 dimensions, i.e. no layered meshes. It is our conviction that in order to be able to resolve all scales important in the ocean circulation, such fully unstructured 3D meshes are a necessity. In this paper a possible solution strategy is explored that provides a similar efficiency for the whole scale of hydrostatic and nonhydrostatic problems without putting any constraints on the type of mesh to be used.

2 Analysis of the equations

Our starting point are the momentum and mass equations for incompressible fluids in the Boussinesq approximation:

$$\frac{D\vec{u}}{Dt} + \frac{\vec{\nabla} p}{\rho_{\text{ref}}} + \vec{\Omega} \times \vec{u} + g\frac{\rho}{\rho_{\text{ref}}}\vec{k} = 0, \tag{1}$$

$$\vec{\nabla} \cdot \vec{u} = 0 \text{ and } \frac{D\rho}{Dt} = 0. \tag{2}$$

Here $\vec{\Omega}$ is the earth's rotation vector and $\vec{k}$ is the vertically upward unit vector. In the following we will use a local coordinate system in which x and y are horizontal coordinates and z is the vertical coordinate. This is only to simplify the analysis in this section. In the model global Cartesian coordinates x, y, z for the whole sphere are used.

Stratification

Typical for ocean flows is a certain amount of stratification in the density function ρ. We can decompose the density into a constant reference density ρ_{ref}, used in the Boussinesq approximation, a static vertically stratified density $\rho_v(z)$ and a dynamic $\delta\rho(x, y, z, t)$:

$$\rho = \rho_{\text{ref}} + \rho_v(z) + \delta\rho(x, y, z, t) \tag{3}$$

Similarly the pressure can be decomposed in a static p_0 of which the vertical gradient is in balance with the ambient stratified density $\rho_{\text{ref}}+\rho_0$, and a dynamic pressure δp:

$$p = p_0(z) + \delta p(x, y, z, t), \quad \frac{\partial p_0}{\partial z} + g(\rho_{\text{ref}} + \rho_0) = 0 \tag{4}$$

This way the static parts of both the buoyancy term and the pressure gradient cancel in the vertical momentum balance, leaving a vertical acceleration term, a dynamic pressure gradient and buoyancy term, and a term for the vertical component of Coriolis.

$$\frac{Dw}{Dt} + \frac{1}{\rho_{\text{ref}}}\frac{\partial \delta p}{\partial z} + g\frac{\delta\rho}{\rho_{\text{ref}}} + \Omega u \cos\phi = 0, \tag{5}$$

where ϕ and u in the vertical Coriolis term are the latitude and longitudinal velocity component respectively.

Stratification and the hydrostatic approximation

In large parts of the ocean the vertical component of the momentum equation is dominated by the hydrostatic balance between the vertical pressure gradient and the buoyancy term $g\delta\rho/\rho_{\text{ref}}$, and the Dw/Dt term may be neglected. An

indication of the relative strengths of Dw/Dt and the buoyancy term, can be found in

$$\frac{D\rho}{Dt} = \frac{D\delta\rho}{Dt} + \frac{\partial\rho_0}{\partial z}w = 0 \tag{6}$$

With $N^2 = -(g/\rho_{\text{ref}})(\partial\rho/\partial z)$, the Brunt-Väisälä frequency, this can be written as:

$$N^2 w = \frac{g}{\rho_{\text{ref}}}\frac{D\delta\rho}{Dt} \tag{7}$$

Then, if T is a typical time scale of the phenomenon under study, a typical vertical velocity is found to be

$$w \approx \frac{g\,\delta\rho}{N^2 T \rho_{\text{ref}}} \tag{8}$$

Therefore for Dw/Dt to be significantly small with respect to the buoyancy term we need:

$$\frac{1}{N^2T^2} \leq 1 \text{ or } T >> N^{-1} \tag{9}$$

In many oceanographic phenomena a typical time scale is given by L/U, with L a typical horizontal length scale and U a typical horizontal velocity. In large parts of the oceans this is significantly higher than the buoyancy time-scale N^{-1}. However, for smaller scale phenomena, or in weakly stratified conditions this is no longer true.

Similarly it can be shown that the vertical component of Coriolis is usually much smaller than the hydrostatic balance terms. In the so called traditional approximation both this term and the vertical acceleration are neglected, leading to a purely hydrostatic balance in the vertical. This is accompanied by a neglect of the $\Omega w \cos\phi$ part of the horizontal Coriolis term, in order to maintain an angular momentum balance.

Decomposition of the pressure

The hydrostatic pressure assumption of the traditional approximation is expressed by

$$\frac{\partial p_{\text{HY}}}{\partial z} + g\rho = 0. \tag{10}$$

Here and in the following we include both the static and the dynamic parts of the pressure. With a boundary condition of $p = 0$ at the surface, the equation is easily solved integrating downward from the surface. In a rigid lid approximation of the ocean however the pressure is allowed to vary freely at the surface. This freedom can be incorporated by adding an arbitrary function that only depends on the horizontal coordinates x and y and thus does not interfere with the hydrostatic balance (10).

$$p(x, y, z) = p_{\text{HY}}(x, y, z) + p_s(x, y) \tag{11}$$

An equation for this surface pressure p_s can be derived by taking the divergence of the momentum equation. For now we will neglect the advective terms, so

$D\vec{u}/Dt = \partial\vec{u}/\partial t$. Thus using the fact that the velocity field is divergence-free, we get:

$$\nabla^2 p + \rho_{\text{ref}}\vec{\nabla}\cdot\left(\vec{\Omega}\times\vec{u}\right) + g\vec{k}\cdot\vec{\nabla}\rho + g\rho\vec{\nabla}\cdot\vec{k} = 0 \tag{12}$$

The divergence of $\vec{k}$ is of order $1/r$, where r is the radius of the earth, and may be neglected. If we assume the pressure to be decomposed as in (11) this reduces to a horizontal PDE for p

$$\nabla_h^2 p + \frac{\partial^2 p}{\partial z^2} + \rho_{\text{ref}}\vec{\nabla}_h\cdot\left(\vec{\Omega}\times\vec{u}\right)_h + \rho_{\text{ref}}\frac{\partial}{\partial z}\left(\vec{\Omega}\times\vec{u}\right)_z + g\frac{\partial\rho}{\partial z} = \tag{13}$$

$$\nabla_h^2 p + \rho_{\text{ref}}\vec{\nabla}_h\cdot\left(\vec{\Omega}\times\vec{u}\right)_h = 0, \tag{14}$$

where subscript h means only the horizontal components of the specified vector. Given the hydrostatic pressure p_{HY} according to (10) with $p_{\text{HY}} = 0$ at the surface, this leaves a horizontal Poisson equation for p_s where the horizontal divergence of the Coriolis should be nearly constant over the vertical for our assumption (11) to be valid.

If the vertical acceleration and Coriolis terms may not be neglected a third non-hydrostatic contribution must be added to the pressure field

$$p(x,y,z) = p_{\text{HY}}(x,y,z) + p_s(x,y) + p_{\text{NH}}(x,y,z) \tag{15}$$

such that the total pressure satisfies the full 3D Poisson equation(12). Solving this equation in combination with the velocity field, we get a solution for the full set of equations with no hydrostatic approximation in the vertical.

3 The solution method for the pressure field

Our aim is to develop a solution strategy in which the full non-hydrostatic equations are solved for, but that is still relatively efficient for large scale problems that are nearly hydrostatic. Like in many implicit non-hydrostatic models, the discretised equations are solved by first calculating a preliminary velocity field out of the discretised momentum equations using a time-stepping method. This new velocity field will not necessarily observe the incompressibility equation $\vec{\nabla}\cdot\vec{u} = 0$. After the new pressure field has been calculated by solving the 3D Poisson equation (12), a correction to the velocity field can be derived such that mass conservation for the incompressible fluid is restored .

This method performs reasonably well for small scale non-hydrostatic problems. The difficulty of efficiently solving larger scale problems with non-hydrostatic models, is that the solution of the discretised 3D Poisson problem is hampered by the variety of length scales. Meshes for such typically large aspect ratio problems are characterised by a very large difference in horizontal and vertical grid distances. This results in a very ill-conditioned matrix associated with the discretised 3D Poisson problem.

The matrix can be split into a horizontal part, the discretisation of $(\partial/\partial x)^2 + (\partial/\partial y)^2$, and a part that is the discretisation of $(\partial/\partial z)^2$. The coefficients of the

horizontal matrix have a typical size of $1/\Delta x^2$ (Δx being a typical horizontal grid length), and those of the vertical matrix will be of size $1/\Delta z^2$ (with Δz being a typical vertical grid size). For large scale ocean models, where $\Delta x >> \Delta z$, this means the vertical coefficients are much larger than those of the horizontal contribution, and specifically the diagonal will be dominated by the $1/\Delta z^2$ contribution.

Because of the difference in horizontal and vertical coefficients, the horizontal coefficients will be neglected in constructing an approximate inverse M^{-1} of the matrix with standard preconditioning techniques for iterative solvers. Horizontal modes in the error $e^k = x_{solution} - x^k$ will therefore be neglected. So for a general iterative method (Krylov subspace methods a.o.), where new search directions are found by applying this approximate inverse to the residual

$$p^k = M^{-1} r^k = M^{-1} \left(b - A x^k \right) = M^{-1} A e^k \tag{16}$$

the iterative method will mainly work in the vertical.

Reduction to horizontal system by vertical lumping

To pick out the horizontal modes out of the residual, the equation should be integrated over the vertical, similarly to the way the surface pressure equation was formed in the previous section. For layered meshes this is done by simply summing over the elements on top of each other. Let P be the operator from the surface mesh to the full 3D mesh, that maps the values of surface nodes to the nodes below. The transpose P^T of this operator, from the full 3D mesh to the surface mesh, then gives the summation in each surface node of all the values in the nodes below. The lumped equation is now given by:

$$P^T A P \tilde{x} = P^T b \tag{17}$$

where $\tilde{x}$ is the solution given in the surface nodes. The correction to the full vector x^k for the horizontal modes in the residual $r^k = b - A x^K$ is then

$$\Delta x_h = P \tilde{x} = P \left(P^T A P \right)^{-1} P^T r^k. \tag{18}$$

Relaxation in the vertical

Adding this contribution, that is completely constant over the vertical, would however introduce a large error in the vertical direction. A relaxation step is needed that reshapes the contribution in the vertical such that the error in this direction is again reduced. Many of the traditional preconditioners could be applied here, since they all represent the vertical component of the matrix well.

An optimal choice in this context can be found using a block approach. All elements in one column are grouped, and the matrix is reordered such that all connections within each column form a diagonal block. Those diagonal blocks

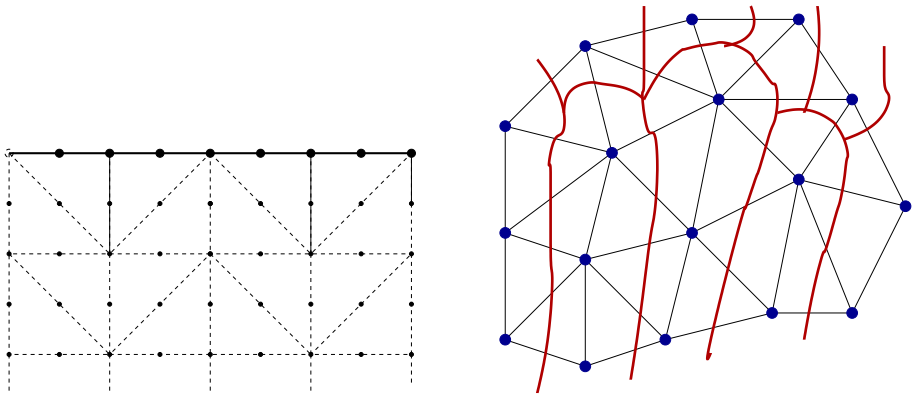

Figure 1: Left: the equations for the full 3D grid are vertically lumped to the surface mesh. Right: in the smoothed aggregation method the nodes in an unstructured meshes are grouped in vertical aggregates.

can be easily inverted, giving a good approximation of the inverse of $(\partial/\partial z)^2$. An optimal relaxation (at least on non-parallel systems) is then given by a Block Forward and Backward Gauss-Seidel (BFBGS) preconditioner step.

After the horizontal contribution Δx_h is subtracted from the residual, the relaxation term with preconditioner M reads

$$\Delta x_r = M^{-1}\left(r - A\Delta x_h\right). \tag{19}$$

In combination with a basic stationary (Richardson-type) iteration, we arrive at the following algorithm

$$\begin{aligned} r^k &= b - Ax^k, \\ \Delta x_h^k &= P\left(P^T AP\right)^{-1} P^T r^k, \\ \Delta x_r^k &= M^{-1}\left(r - A\Delta x_h\right), \\ x^{k+1} &= x^k + \Delta x_h^k + \Delta x_r^k \end{aligned} \tag{20}$$

To enhance the solution of any non-hydrostatic parts this can be further accelerated in combination with a conjugate gradient iteration. The combination of a horizontal solve with a vertical block preconditioner is similar to what can be found in [1].

4 Lumping over the vertical in multigrid context

The solution approach in the previous section is based on one important assumption, namely that the mesh is layered in the vertical. The importance of

fully unstructured 3D meshes in ocean modelling was already outlined in the introduction. The problem of a variety of length scales, and the need to to efficiently solve both the small and large scale flow features only becomes more urgent in this case. Furthermore it is not directly clear how an equivalent for the vertical lumping can be constructed for unstructured meshes.

A step forward can be made by realising that the iteration between the 2D horizontal pressure solve and the fully 3D solve can be seen as a 2-level multigrid cycle. In this context the fine grid consists of the complete 3D mesh whereas the coarser grid is given by the surface mesh. Our algorithm contains all the ingredients that make up a basic multigrid approach

- A selection of the coarse grid. In the layered case this is the surface mesh
- A restriction operator P^T from the fine to the coarse mesh and a prolongator P from the coarse to the fine mesh.
- The construction of a coarse grid operator, usually given by the Galerkin operator $A_c = P^T AP$.
- A smoothing operation to reduce the smooth error after going from the coarse to the fine grid.

Algebraic multigrid

Classical multigrid approaches base the selection of the coarse mesh on a coarsening operation on the associated computational mesh, for instance $h \rightarrow 2h$ coarsening, or the collapsing of vertical columns in the previous section. Algebraic multigrid on the other hand bases this selection purely on the algebraic structure of the matrix. This makes them applicable to a much wider range of problems, such as those on unstructured meshes.

A variety of different multigrid approaches exists, differing in how exactly they select the coarse grids and how they construct the interpolation and restriction operators.

Some approaches derive this purely from the topological structure of the matrix connectivity graph. For our application however it is important that also the magnitude of the coefficients is taken into account. This makes it possible to distinguish between horizontal and vertical connections. Because as we saw the vertical coefficients are much larger, this means the coarsening strategy will focus on those connections and the coarsening will mainly take place in that direction.

An example is the *smoothed aggregation* approach[2], in which nodes that are strongly connected are grouped together. Those groups (or aggregates) form the nodes of the coarse mesh. In the layered case this automatically leads to column-like groups of nodes, but also for unstructured meshes the groups will be mainly vertical. Another important component of this approach is the fact that smoothing is already introduced in the prolongation and restriction operators. First a basic tentative prolongation operator P is constructed that simply copies

the value of the coarse node to all the nodes of the associated aggregate in the fine mesh. The actual prolongator is then given by this operation followed by a smoothing step:

$$\tilde{P} = SP, \text{ and similarly } \tilde{P}^T = P^T S, \tag{21}$$

where the smoother S needs to be symmetric. This means that the smoothing is also incorporated in the construction of the coarse grid operator $\tilde{P}^T A \tilde{P}$.

5 Summary and outlook

The application of algebraic multigrid techniques described in this paper promises to provide an efficient solution strategy for large scale ocean modelling with unstructured meshes, that performs optimally for both hydrostatic and non-hydrostatic problems, and is able to vary smoothly its behaviour e.g. going from the basin scale to the convective scale.

Finally a further possible application of algebraic multigrid techniques could be found by trying to directly solve the flow equations, both the momentum and mass equations, in a coupled way. This would ask for even more of the flexibility of algebraic multigrid to deal with various scales.

References

[1] J. Marshall, A. Adcroft, C. Hill, L. Perelman, and C. Heisey. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. *Journal of Geophysical Research*, 102:5753–5766, 1997.

[2] P. Vanek, J. Mandel, and M. Brezina. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing*, 56, 1996.